Exhibit 5

U.S. Patent No. 8,018,852 ("'852 Patent")

Exemplary Accused Product

NetApp products, including at least each of the following products (and their variations) infringe at least Claim 1 of the '852 Patent: NetApp ONTAP OS and storage and network controllers running ONTAP OS. The infringement chart below is based on the NetApp ONTAP 9 OS ("ONTAP 9"), which is exemplary of the infringement of the '852 Patent.

Claims	ONTAP 9
[1pre]. A equal-cost source-resolved (ECSR) port selection method for transmitting an outbound flow, for a conversation, from a first node to a second node in a distributed network comprising one or more routing devices, the method comprising the steps of:	ONTAP 9 enabled devices perform an equal-cost source-resolved (ECSR) port selection method for transmitting an outbound flow from a source node to a destination node in a distributed network comprising one or more routers. **Port-based load balancing** You can equalize traffic on a multimode interface group based on the transport layer (TCP/UDP) ports by using the port-based load-balancing method. The port-based load-balancing method uses a fast hashing algorithm on the source and destination IP addresses along with the transport layer port number. https://library.netapp.com/ecmdocs/ECMP1114171/html/GUID-44BF66B6-A777-4152-B4D3-8ECB221EBBD9.html 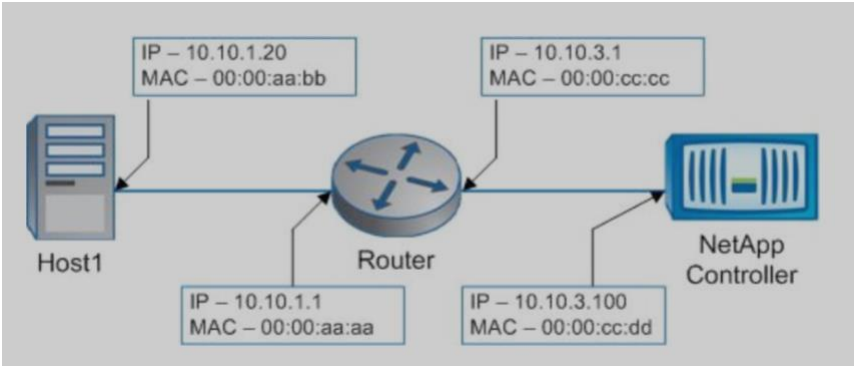

	https://www.netapp.com/us/media/tr-3802.pdf 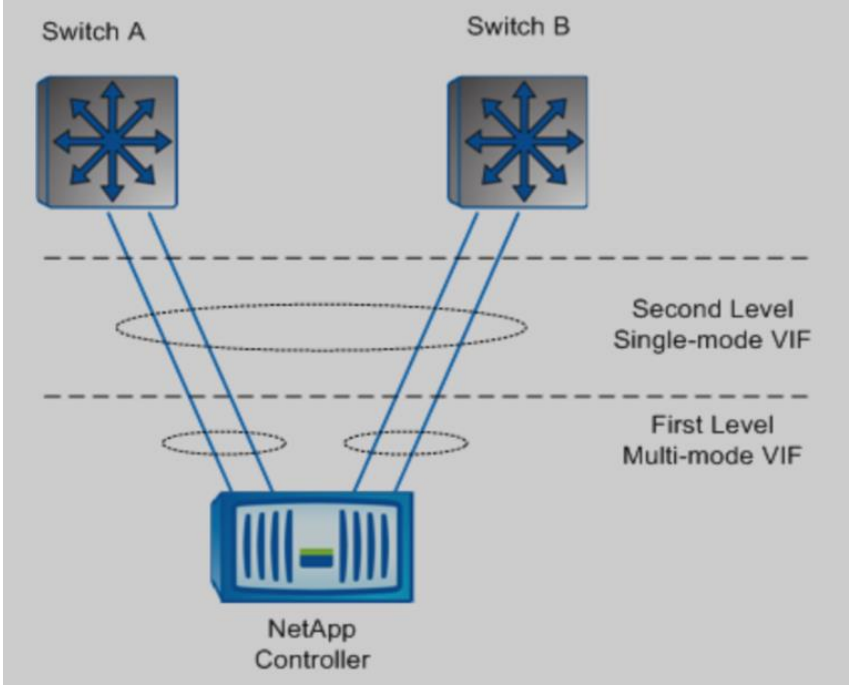 https://www.netapp.com/us/media/tr-3802.pdf
[1a] at a switching device:	ONTAP 9 is executed on a switching device.

	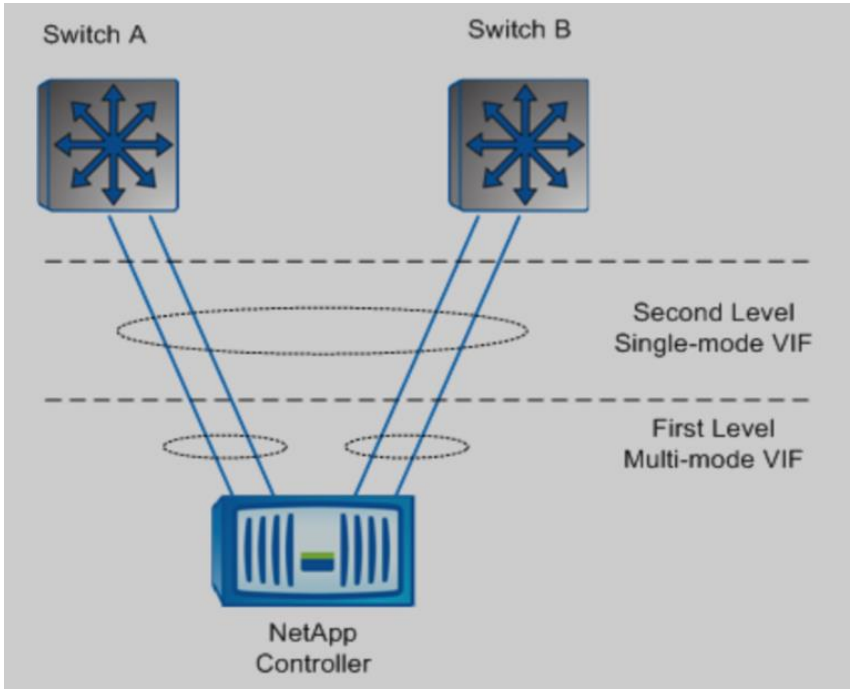 https://www.netapp.com/us/media/tr-3802.pdf
[1b] identifying a plurality of ports associated with minimal equal-cost paths from the first node to the second node;	ONTAP 9 identifies a plurality of ports associated with minimal equal-cost paths from the first node to the second node. • Ports: – Physical ports: These ports are used for various functions and can have different types of configurations. – Logical ports: Virtual LANs (VLANs)/(802.1Q) and interface groups (IFGRPs) make up the options for logical ports. https://www.netapp.com/us/media/tr-4182.pdf

	Port bonding methods are available on both the network and storage components to aggregate multiple physical links into a single virtual interface. A load balancing algorithm distributes traffic across the physical links in the channel to provide additional performance and redundancy within the network. The terminology for this technology sometimes differs between NetApp and the networking industry, which can cause confusion. While the network industry uses terms such as EtherChannel or port-channel, NetApp calls these Multi-Mode VIFs, or virtual interfaces. An incorrect term that's sometimes used is "trunked interfaces" or "trunk". The term "trunk" or "trunking" should be avoided, since the networking industry already uses the term trunk to describe VLAN trunking. **Load balancing in multimode interface groups** You can ensure that all interfaces of a multimode interface group are equally utilized for outgoing traffic by using the IP address, MAC address, round-robin, or port based load-balancing methods to distribute network traffic equally over the network ports of a multimode interface group. https://library.netapp.com/ecmdocs/ECMP1114171/html/GUID-27EAC723-A07A-4F60-A1AC-8E91AE6AE82E.html
[1c] if an inbound flow, for the conversation, received from the second node is detected on a first port of the plurality of ports of the first node, then associating a network address of the inbound flow with the first port on which it was received; and	ONTAP 9 associates a network address of the inbound flow with the first port on which it was received if an inbound flow, for the conversation, received from the second node is detected on a first port of the plurality of ports of the first node. Fast path is an alternative routing mechanism to the routing table, in which the responses to incoming network traffic are sent back by using the same interface as the incoming traffic. By avoiding the routing table lookup, fast path provides a shorter route for data access. Fast path is used in all TCP and NFS/UDP protocols. By default, fast path is enabled on the storage system; however, you can disable it. https://kb.netapp.com/app/answers/answer_view/a_id/1002602/~/how-does-ip-fast-path-function- **Load balancing in multimode interface groups** You can ensure that all interfaces of a multimode interface group are equally utilized for outgoing traffic by using the IP address, MAC address, round-robin, or port based load-balancing methods to distribute network traffic equally over the network ports of a multimode interface group.

	https://library.netapp.com/ecmdocs/ECMP1114171/html/GUID-27EAC723-A07A-4F60-A1AC-8E91AE6AE82E.html **Port-based load balancing** You can equalize traffic on a multimode interface group based on the transport layer (TCP/UDP) ports by using the port-based load-balancing method. The port-based load-balancing method uses a fast hashing algorithm on the source and destination IP addresses along with the transport layer port number. https://library.netapp.com/ecmdocs/ECMP1114171/html/GUID-44BF66B6-A777-4152-B4D3-8ECB221EBBD9.html 4. Port - Traffic balancing based on source and destination IP addresses along with the transport layer port number. Also referred to as IP+port. Available in Data ONTAP versions 7.3.1D3 and post 7.3.2. Because this choice adds the TCP/UDP port number to the calculation, a single host can have multiple TCP streams to the Storage Controller, each stream using a different member interface within the VIF/ifgrp. This can allow for bandwidth usage that is higher than the throughput possible in a single member interface. https://kb.netapp.com/app/answers/answer_view/a_id/1001531/~/how-does-load-balancing-on-a-vif-work%3F- As of Data ONTAP 7.3.2, a multimode or LACP ifgrp/vif uses an implementation of a "SuperFastHash", utilizing the last 16 bits of the source and destination IP addresses (-b ip), the last 16 bits of the source and destination MAC addresses (-b mac), or the last 16 bits of the source and destination IP addresses in combination with the source and destination TCP port (-b port). The output of the algorithm results in a far more dynamic, more balanced distribution than the algorithm used in versions of Data ONTAP prior to 7.3.2. The result is still the same, however, in that each TCP stream will associate with only one interface, allowing for only one port's worth of bandwidth per TCP stream. https://kb.netapp.com/app/answers/answer_view/a_id/1001531/~/how-does-load-balancing-on-a-vif-work%3F-
[1d] transmitting the outbound flow from the first port of the first node to the second node based on the network address associated with the inbound flow.	ONTAP 9 transmits the outbound flow from the first port of the first node to the second node based on the network address associated with the inbound flow.

	How fast path works Fast path is an alternative routing mechanism to the routing table. In fast path, the responses to incoming network traffic are sent back by using the same interface as the incoming traffic. By avoiding the routing table lookup, fast path provides a quick access to data. If fast path is enabled on an interface group and a physical interface in that group receives an incoming request, the same physical interface might not send a response to the request. Instead, any other physical interface in an interface group can send the response. **How fast path works with NFS/UDP** NFS/UDP traffic uses fast path only when sending a reply to a request. The reply packet is sent out on the same network interface that received the request packet. For example, a storage system named toaster uses the toaster-e1 interface to send reply packets in response to NFS/UDP requests received on the toaster-e1 interface. Fast path is used only in NFS/UDP. However, fast path is not used in other UDP-based NFS services such as portmapper, mountd, and nlm. https://library.netapp.com/ecmdocs/ECMP1114171/html/GUID-8276014A-16EB-4902-9EDC-868C5292381B.html Fast path is an alternative routing mechanism to the routing table, in which the responses to incoming network traffic are sent back by using the same interface as the incoming traffic. By avoiding the routing table lookup, fast path provides a shorter route for data access. Fast path is used in all TCP and NFS/UDP protocols. By default, fast path is enabled on the storage system; however, you can disable it. https://kb.netapp.com/app/answers/answer_view/a_id/1002602/~/how-does-ip-fast-path-function- ◦ Load balancing between multiple network interfaces on the same subnet. Load balancing is achieved by sending responses on the same interface of your storage system that receives the incoming requests. If fast path is enabled on an interface group and a physical interface in that group receives an incoming request, fast path does not govern which physical interface group member port sends a response to the request (this is specifically related to LACP or static multimode interface group). In an interface group, the interface selected when transmitting a packet is determined based upon the configured load balancing policy, and incoming traffic will arrive based upon how the neighboring switch utilizes its own load balancing policy to select ports when transmitting traffic. ◦ Improved storage system performance by skipping routing table lookups. For more information on how routing in Data ONTAP works, see 1014265: Triage Template - How to troubleshoot Default Route/Routing Issues in Data ONTAP 7-Mode https://kb.netapp.com/app/answers/answer_view/a_id/1002602/~/how-does-ip-fast-path-function-

	Remember, the load balancing in an IFGRP occurs on outbound traffic, not inbound traffic. So, when a response is being sent back to a requester, the load-balancing algorithm comes into play to determine which "path" is optimal to use to send the response back. Also, it's important to note that the preceding load-balancing options might differ from other settings in the environment. You should thoroughly check other devices that might be connected to the ports on the nodes in the cluster. If you use "port" on the IFGRP configuration on the cluster, make sure that the switch port on the Cisco, Juniper, Brocade, or other device is also configured in the same way. https://www.netapp.com/us/media/tr-4182.pdf 4. Port - Traffic balancing based on source and destination IP addresses along with the transport layer port number. Also referred to as IP+port. Available in Data ONTAP versions 7.3.1D3 and post 7.3.2. Because this choice adds the TCP/UDP port number to the calculation, a single host can have multiple TCP streams to the Storage Controller, each stream using a different member interface within the VIF/ifgrp. This can allow for bandwidth usage that is higher than the throughput possible in a single member interface. https://kb.netapp.com/app/answers/answer_view/a_id/1001531/~/how-does-load-balancing-on-a-vif-work%3F- As of Data ONTAP 7.3.2, a multimode or LACP ifgrp/vif uses an implementation of a "SuperFastHash", utilizing the last 16 bits of the source and destination IP addresses (-b ip), the last 16 bits of the source and destination MAC addresses (-b mac), or the last 16 bits of the source and destination IP addresses in combination with the source and destination TCP port (-b port). The output of the algorithm results in a far more dynamic, more balanced distribution than the algorithm used in versions of Data ONTAP prior to 7.3.2. The result is still the same, however, in that each TCP stream will associate with only one interface, allowing for only one port's worth of bandwidth per TCP stream. https://kb.netapp.com/app/answers/answer_view/a_id/1001531/~/how-does-load-balancing-on-a-vif-work%3F- Note that because the last octet of the source and destination IP address are the only parts that are used in the load balancing calculation, devices that share a common last octet will end up hashing to the same physical link across a channel, even if they are in different subnets. https://www.netapp.com/us/media/tr-3802.pdf